

Please type a plus sign (+) inside this box. → ☐

02/04/00

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

Attorney Docket No.

22490-0701-00US

Total Pages

23

First Named Inventor or Application Identifier

Only for new nonprovisional applications under 37 CFR 1.53(b)

Express Mail Label No.

EL371023895US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents:

Assistant Commissioner for Patents

ADDRESS TO: Box Patent Application
Washington, DC 20231

1. ☒ Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)

2. ☒ Specification [Total Pages 15]
(preferred arrangement set forth below)

- Descriptive title of the Invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to Microfiche Appendix
- Background of the invention
- Brief Summary of the Invention
- Brief Description of the Drawings (if filed)
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

3. ☒ Drawings (35 USC 113) [Total Sheets 6]
(INFORMAL)

4. Oath or Declaration [Total Pages]

a. ☐ Newly executed (original or copy)

b. ☐ Copy from a prior application (37 CFR 1.63(d))
(for continuation/divisional with Box 17 completed)
[Note Box 5 below]

i. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in
the prior application, see 37 CFR 1.63(d)(2) and 1.33(b)

5. ☐ Incorporated By Reference (useable if Box 4b is checked) The entire
disclosure of the prior application, from which a
copy of the oath or declaration is supplied under Box 4b,
is considered as being part of the disclosure of the accompanying
application and is hereby incorporated by reference therein.

6. ☐ Microfiche Computer Program (Appendix)

7. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)

a. ☐ Computer Readable Copy

b. ☐ Paper Copy (identical to computer copy)

c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & document(s))

9. ☐ 37 CFR 3.73(b) Statement ☐ Power of Attorney
(when there is an assignee)

10. ☐ English Translation Document (if applicable)

11. ☐ Information Disclosure ☐ Copies of IDS
Statement (IDS)/PTO-1449 Citations

12. ☐ Preliminary Amendment

13. ☒ Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)

14. ☐ Small Entity ☐ Statement filed in prior application,
Statement(s) Status still proper and desired

15. ☐ Certified Copy of Priority Document (s)
(if foreign priority is claimed)

16. ☐ Other:

17. If a CONTINUING APPLICATION,
check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part
(CIP)

of prior application No: /

18. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Label

(Insert Customer No. or Attach Bar code label here)

or

☒

Correspondence address below

NAME	Alan S. Hodes - McCutchen, Doyle, Brown & Enersen, L.L.P.				
ADDRESS	Three Embarcadero Center				
CITY	San Francisco	STATE	CA	ZIP CODE	94111-4067
COUNTRY	USA	TELEPHONE	415-393-2000	FAX	415-393-2286

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

MULTI-THREADED FRAME SAFE SYNCHRONIZATION OF A SIMULATION

Technical Field

The present invention relates to simulations and, in particular, to a
5 simulation in which the various aspects of the simulation are executed
asynchronously but communicate with each other in a safely synchronized
manner.

Background

Virtual prototyping solutions can greatly accelerate product
10 development. Obviously, the more accurately a solution simulates "reality",
the more confidence one can have in that solution. On the other hand,
though, the most useful simulation is typically the one that not only
simulates reality faithfully, but also takes the fewest resources to do so.

Whereas reality is continuous in nature, computer simulations are
15 forced to approximate the continuity by acting upon and examining and/or
evaluating the simulation model only at discrete points in time. The greater
the frequency at which the simulation model is examined, the more
accurate is the simulation, since continuity is more closely approached.
Thus, conventionally, simulations are designed such that the evaluation rate
20 is high enough to sufficiently approximate reality (i.e., with acceptable error
and without noticeable jumps in display), but not so fast that the
computer's resources are completely exhausted.

The evaluation rate required depends largely upon the rate at which
the simulated events would occur in reality. For example, if modeling the
25 fast-moving parts within an engine, the evaluation rate should be relatively
high. On the other hand, modeling the slow buildup of heat within the

PATENT

engine could be carried out at a much slower rate. Furthermore, where visual interaction is important, evaluation for display should be at a relatively high rate if the sensation of interactivity is to be provided to the user.

5 Since different aspects of a simulation occur with varying evaluation rate requirements, it would be preferable if each aspect executed at a rate most suitable to it. Unfortunately, conventional simulations force the entire simulation to execute at a common frequency. The common frequency is generally faster than required for some aspects of the simulation and
10 perhaps slower than ideally required for other aspects of the simulation. For example, if the calculations of engine heat buildup could only be performed at once or twice a second due to the complexity of the calculations, the calculation of engine part positions could only be performed one or twice a second as well, even though it would be a much
15 simpler calculation. Furthermore, visualization of movement around the engine would be seriously hampered, with position updates at only the once or twice per second rate.

 A typical solution to the evaluation rate problem discussed above is to perform the simulation in batch, allowing the user to inspect the results
20 in detail once the simulation is completed. A disadvantage of this approach, however, is that the simulation cannot be interactive.

Summary

 In accordance with the invention, a computer program is provided to
25 execute a simulation. The computer program includes a plurality of service program means, where the services are configured to collectively determine simulated attributes of objects of an environment under simulated

operation. Each service is associated with at least a subset of object attributes in an object database. Furthermore, each service is executing at a rate independent of the other services, and at least some of the services access and operate upon object attributes, from the object database, with which the service is associated.

A write queue program is associated with each service, to queue write requests from the service to write determined simulated attributes to the object database. Nodes coordinate execution of the queued requests to cause the determined simulated attributes to be written to the object database in a manner such that each service has a coherent view of all the object attributes.

Brief Description of the Figures

Fig. 1 broadly illustrates a simulation environment 100 in accordance with an embodiment of the invention.

Fig. 2 illustrates processing within a service of the Fig. 1 environment 100.

Fig. 3 illustrates how a node change request is handled by a node.

Fig. 4 illustrates how a node processes a request from a service to ensure the service is pointing to the most recently determined value for an object (or portion of an object).

Fig. 5 is a time line that broadly illustrates how services interact.

Fig. 6 is a time line that illustrates the interaction between services in greater detail.

Detailed Description

A simulation environment 100 is now described, first broadly with reference to Fig. 1. Block 102 is an object database. The "attributes" of "objects" in the database 102 represent the "state" of the environment
5 being simulated. Block 104 illustrates the execution of "services" (in this case, services A through C) that constitute the active (computational) part of the simulation.

Block 106 illustrates a "data probe context". Input and output data probes are computational objects (implemented as services in the service
10 execution block 104) that provide the simulation user an interface to attributes of objects in the object context 102. Input data probes inject values from some source (e.g., a file or a connection to another application) into attributes of an object during a simulation. Conversely, output data probes examine the values of attributes during a simulation and provide
15 those values to some destination such as a file or a connection to another application. Data probes are described later in greater detail.

Finally, block 108 illustrates a "behavior context". Behavior is the programming employed by a service (in the service execution block 104) that defines how a service is to affect the modeled environment during a
20 simulation. One type of behavior is "logic", which is an event-based state machine description of behavior. Another type of behavior is "mechanism", which is a physics, kinematics-based description of behavior.

Now, the operation of the services execution block 104 is described in greater detail. As alluded to above, services are objects that perform
25 calculations or other functionality repeatedly to accomplish the simulation. Typically, multiple services exist simultaneously and the services each perform their functionality at a particular temporal frequency. Theoretically, the number of services executing at one time is unlimited,

PATENT

although as a practical matter the number of services may be limited by the operating system or by computational performance limitations.

In accordance with an embodiment of the invention, the multiple services operate asynchronously, but data communication among the
5 multiple asynchronously operating services is safely synchronized. This is accomplished even while accommodating the various asynchronous "frame rates" of the services.

To support the multiple "views" of object data required by various services. All objects to which access is required by a service are derived
10 from a "node" base class. The node interacts with services to provide frame safe access to its data members. Each node may maintain multiple "node images" for an object (or for a portion of an object). Each service potentially sees different sets of changes at various times, although all services are ultimately working on the same values. In one embodiment, to
15 minimize the impact on memory, the data duplication occurs at a relatively fine granularity, within objects referred to as "nodes". Nodes are typically relatively small units (portions of object), and data copies are maintained only for those nodes requiring them. Thus, memory duplication is minimized.

20 In one embodiment, a node does not actually contain attribute data itself. Rather, a node serves as a wrapper to one or more node images that contain the attribute data. While a node image actually contains data, a node includes "frame-safe" methodology to provide the services access the data in the node images. That is, in essence, a node hides the "multi-
25 service issues" from the services. A service that needs to access the value of an object attribute requests the value from or writes the value to, a node corresponding to that object (or to the portion of the object that includes the attribute).

PATENT

The service passes an ID of the service as an argument to the node. The node uses the ID internally to manage multiple node images it may contain. The node determines which node image is accurate for the requesting service, and creates and deletes node images as needed to
5 maintained the required "clean" copies.

Before discussing in greater detail how a service accesses node data, though, it is useful to briefly discuss the "phases" of processing that a service executes. Fig. 2 illustrates the processing phases of a typical service. First, the service performs its computations in a computational
10 phase 202. During this computational phase 202, the service queues up (but does not yet process) "change requests". A "change request" is a request to a node to change an attribute of an object represented by the node. Also, during the computational phase 202, the service requests reading of attribute values required for its processing. The requests to read
15 attribute values are not queued but, rather, are processed as part of the computational phase 202.

Next, the service enters a change request processing phase 204 in which it actually processes (or causes to be processed) the change requests queued up during the computational phase 202. In one embodiment, the
20 change request processing by the service consists simply of looping through the queued change requests and passing them to the appropriate nodes for execution. Alternately, the service could notify the nodes (or some other program) that change requests are in its queue for processing.

Finally, if the time period for this frame of the service has not
25 expired by the end of the change request processing phase 204, the service enters an idle phase 206 during which it does nothing substantive until it is time to enter the computational phase 202 again.

PATENT

The processing of change requests is now described in greater detail. Fig. 3 illustrates the processing performed by a node in response to a change request passed from a service. After starting (302), at step 304 it is determined if the node image corresponding to the service is the "latest image" or if another service is referencing this node. What is meant by "latest image" is that this node image represents the latest updated state of the object represented by the node. (The process by which a node image becomes the "latest image" is described later.)

If either the node image is the latest image or a service other than the requesting service is referencing the node, then at step 306 a new node image is created. Also, the node records that the latest updated state of the object represented by the node (i.e., "latest image") is this new node image. The old node image is deleted.

Then (or if the determination at step 304 is "NO"), at step 308 other services for which the node is maintaining node images are notified that the "latest image" has changed. At step 310, the requested write operation is made to the node image for the service (whether the "old" node image corresponding to the service or the newly created node image corresponding to the service).

The notification (step 308), and how the notified services react to a notification, is now described. Referring back to Fig. 2, a notification is processed by a service as a "null" change request during the change request processing step 204. In particular, in response to such a notification, the notified service creates a change request for the node for which has been changed by the notifying service. The change request is "null", though, because it does not include data to be written to the object represented by the node. Rather, the change request is a request to make

the node image associated with the notified service corresponds to the "latest image" if required.

Fig. 4 illustrates how a node processes a "null" change request. First, the processing starts at step 402. At step 404, a determination is made as to whether the node image presently associated with the calling service (i.e., the notified service) is the "latest" node image. If not, then there is no further processing of the "null" change request. If the node image presently associated with the calling service is not the "latest" node image, then the node associates the calling service with the "latest" node image, and then deletes the "old" node image for the calling service.

Having described the operation of a service in isolation and also how a node reacts to change requests by a service, reference is now made to Figs. 5 and 6 to describe how services executing in accordance with the previous discussion coexist in apparent synchronicity. Fig. 5 is a timeline that shows how each service executes in its own "frame time". In Fig. 5, the actual time a particular service spends in a frame differs from the time spent in a frame by another service, although this isn't shown in Fig. 5. That is, the services are not synchronized as would appear from Fig. 5 on its face. In fact, though, as is shown later with reference to Fig. 6, it appears to each service that it is in fact operating synchronously with the other services.

Turning now to Fig. 5, and looking at Frame N first, each service reads object attributes from an object "frame context" 502 that appears to that service to be synchronous with the frame execution of the service. (Remembering the previous discussion, this is accomplished in one embodiment by the service generating a read request to a node object, and the node object providing the requested attribute data from a "latest" node image associated with the node.) During its computational phase (202,

also shown in Fig. 2), the service queues "change requests" (indicated in Fig. 5 as WriteOp's) to write object attribute data back to the object context 502. After the computational phase, each service processes its queued change requests during a protected critical section by asking the

5 appropriate nodes to perform the changes indicated by the change requests (as shown in Fig. 3). In one embodiment, during the "critical section" for processing a change to a node, the node may not otherwise be accessed, either for read access or for write access. Fig. 5 further illustrates similar processing occurring in Frame N+1.

10 Now, Fig. 6 is a timeline in "real time" that illustrates how the node processing provides the appearance of synchronization to the services. Taking service A first, it can be seen (as denoted by arrow 602a) that Service A is provided a node attribute (in shorthand, "reads" the attribute, even though such "reading" is via the node object) for processing. At some

15 point later (denoted by arrow 604a), Service A queues a write operation of an object attribute. Yet later (denoted by arrow 606a), during a protected "critical" section of processing, Service A causes (requests) the appropriate node to write the object attribute to an image (the current – or "latest" – image). This is similarly occurring for Service B and Service C, shown in

20 Fig. 6 using the same reference numerals as for Service A but with subscripts B and C, respectively.

That is, Service B reads object attributes (602b); computes on the object attributes and queues write operations of object attributes (604b); and causes the object attributes of the queued write operations to be

25 associated with the current image for the node (606b). A similar processing occurs for Service C (602c, 604c and 606c).

It should be noted that the invention is to be defined by the claims, and not by the embodiment(s) just described.

WHAT IS CLAIMED IS:

- 1 1. A computer program product for use with a computer system to execute
2 a simulation, comprising:
3 a plurality of service computer readable program code means,
4 the service program code means configured to collectively determine
5 simulated attributes of objects of an environment under simulated
6 operation, each service program code means associated with at least a
7 subset of object attributes in an object database and each service program
8 code means executing at a rate independent of the other service program
9 code means, at least some of the service program code means including
10 computer readable program code means to access and operate upon object
11 attributes, from the object database, with which the service program code
12 means is associated;
13 write queue computer readable program code means
14 associated with each service program code means that queues write
15 requests from the service program code means to write determined
16 simulated attributes to the object database; and
17 node computer readable program code means that
18 coordinates execution of the queued requests to cause the determined
19 simulated attributes to be written to the object database in a manner such
20 that each service program code means has a coherent view of all the object
21 attributes.
- 1 2. The program product of claim 1, wherein
2 the node program code means includes computer readable
3 program code means for creating an image of at least a portion of an object

PATENT

4 whose attribute is to be written to the object database and for writing the
5 determined simulated attributes to the image; and
6 to write the determined simulated attributes of the object to
7 the object database, the node program code means associates the image
8 with the object database.

1 3. The program product of claim 2, wherein:
2 the node program code means associates the image with the
3 object database by changing a pointer for the object in the object database
4 to point to the image.

1 4. The program product of claim 2, wherein:
2 the node code program means includes computer readable
3 program code means for notifying at least some of the service program
4 means that the node program code means is associating an image with the
5 object database.

1 5. The program product of claim 4, wherein:
2 in response to a service program code means receiving an
3 object database association notification from a node program code means,
4 the write queue program code means associated with that service program
5 code means queues a request to the node program means to synchronize
6 that service program code means to the image, and
7 the node program code means includes computer readable
8 program code means for synchronizing the service to the image.

1 6. The program product of claim 1, wherein:
2 the service program code means each include a computational
3 phase during which it operates upon the object attributes and during which
4 the write queue program code means queues the write requests generated
5 by the service program code means during the computational phase; and
6 the write requests queued for a particular service program
7 code means are processed by the node program code means during a write
8 request processing phase that is outside the computational phase.

1 7. The program product of claim 4, wherein the at least some of the
2 service program code means which the node program code means notifies
3 includes service program code means that are associated with the object
4 attributes represented by the node image.

1 8. A method executed by a computer to accomplish a simulation,
2 comprising:
3 a plurality of service steps that collectively determine
4 simulated attributes of objects of an environment under simulated
5 operation, each service step associated with at least a subset of object
6 attributes in an object database and each service step executing at a rate
7 independent of the other service steps, at least some of the service steps
8 including steps to access and operate upon object attributes, from the
9 object database, with which the service step is associated;
10 a write queue computer step associated with each service
11 step that queues write requests from the service step to write determined
12 simulated attributes to the object database; and
13 a node step that coordinates execution of the queued
14 requests to cause the determined simulated attributes to be written to the

15 object database in a manner such that each service step has a coherent
16 view of all the object attributes.

1 9. The method of claim 8, wherein

2 the node step includes a step of creating an image of at least
3 a portion of an object whose attribute is to be written to the object
4 database and of writing the determined simulated attributes to the image;
5 and

6 to write the determined simulated attributes of the object to
7 the object database, the node step associates the image with the object
8 database.

1 10. The method of claim 9, wherein:

2 the node step associates the image with the object database
3 by changing a pointer for the object in the object database to point to the
4 image.

1 11. The method of claim 9, wherein:

2 the node step includes a step of notifying at least some of the
3 service steps that the node step is associating an image with the object
4 database.

1 12. The method of claim 11, wherein:

2 in response to a service step receiving an object database
3 association notification from a node step, the write queue step associated
4 with that service step queues a request to the node step to synchronize
5 that service step to the image, and

6 the node step includes a step of synchronizing the service to
7 the image.

1 13. The method of claim 8, wherein:

2 the service steps each include a computational phase during
3 which it operates upon the object attributes and during which the write
4 queue step queues the write requests generated by the service step during
5 the computational phase; and

6 the write requests queued for a particular service step are
7 processed by the node step during a write request processing step that is
8 outside the computational step.

1 14. The method of claim 11, wherein the at least some of the service steps
2 which the node step notifies includes service steps that are associated with
3 the object attributes represented by the node image.

ABSTRACT

A computer program executes a simulation. The computer program includes a plurality of service programs, where the service programs are configured to collectively determine simulated attributes of objects of an environment under simulated operation. Each service is associated with at least a subset of object attributes in an object database. Furthermore, each service is executing at a rate independent of the other services, and at least some of the services access and operate upon object attributes, from the object database, with which the service is associated. A write queue program is associated with each service, to queue write requests from the service to write determined simulated attributes to the object database. Nodes coordinate execution of the queued requests to cause the determined simulated attributes to be written to the object database in a manner such that each service has a coherent view of all the object attributes.




Fig. 1

Fig. 1 is a block diagram of a system 100. The system 100 includes a behavior context 108, a data probe context 106, an object context 102, and a service execution 104. The behavior context 108 includes an attribute changer B and an attribute changer A. The data probe context 106 is connected to the object context 102 and the service execution 104. The object context 102 is connected to the attribute changer B and the attribute changer A. The service execution 104 includes service A, service B, and service C. The attribute changer A is connected to service A, service B, and service C. The attribute changer B is connected to the object context 102. The data probe context 106 is connected to the object context 102 and the service execution 104.

Fig. 2

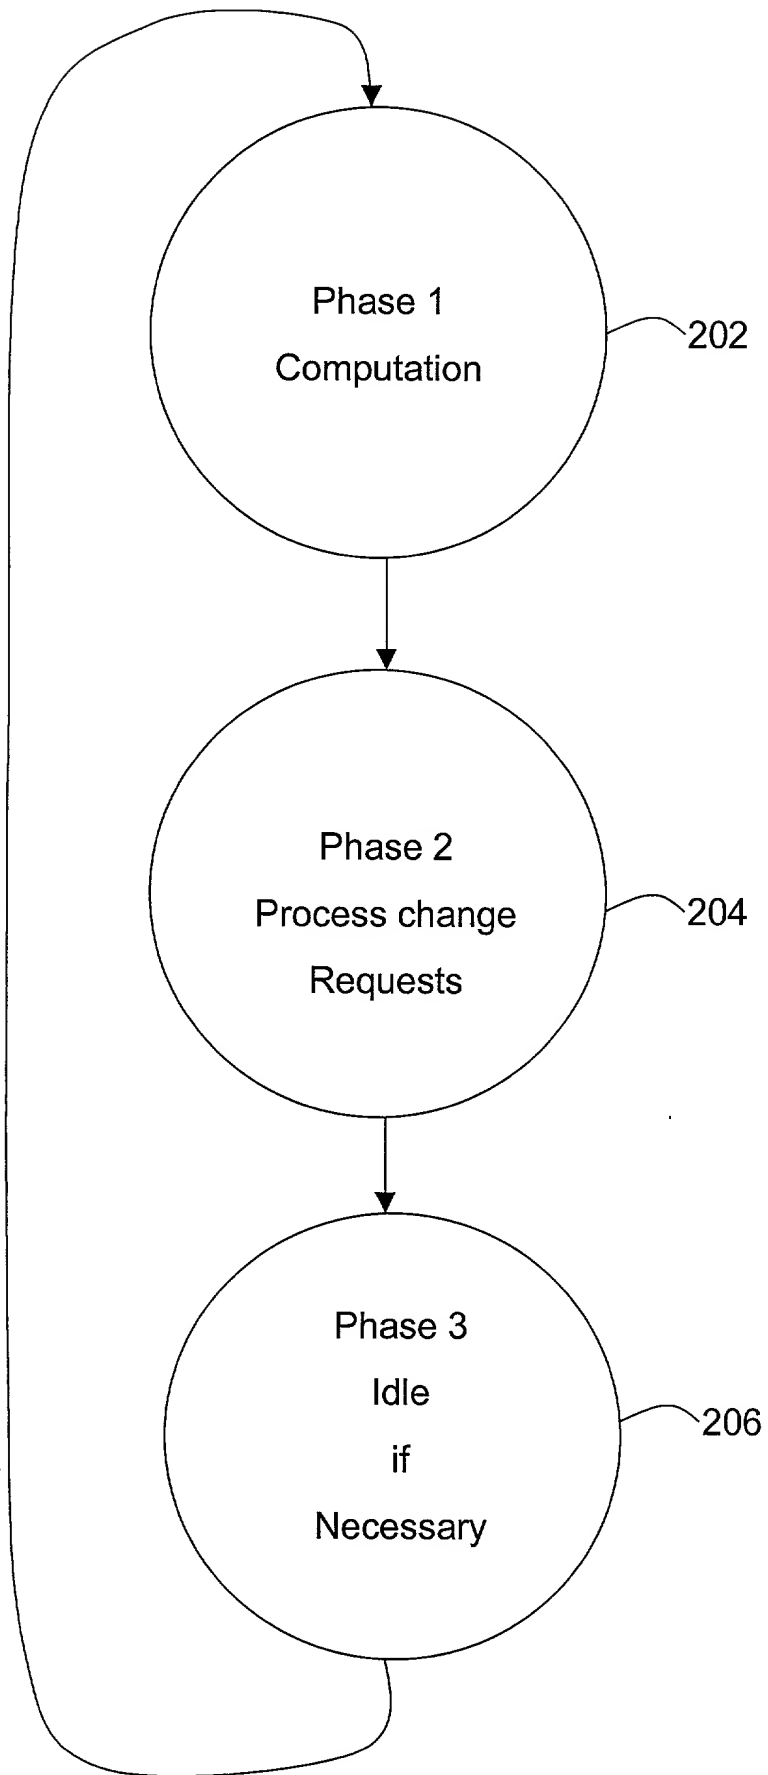
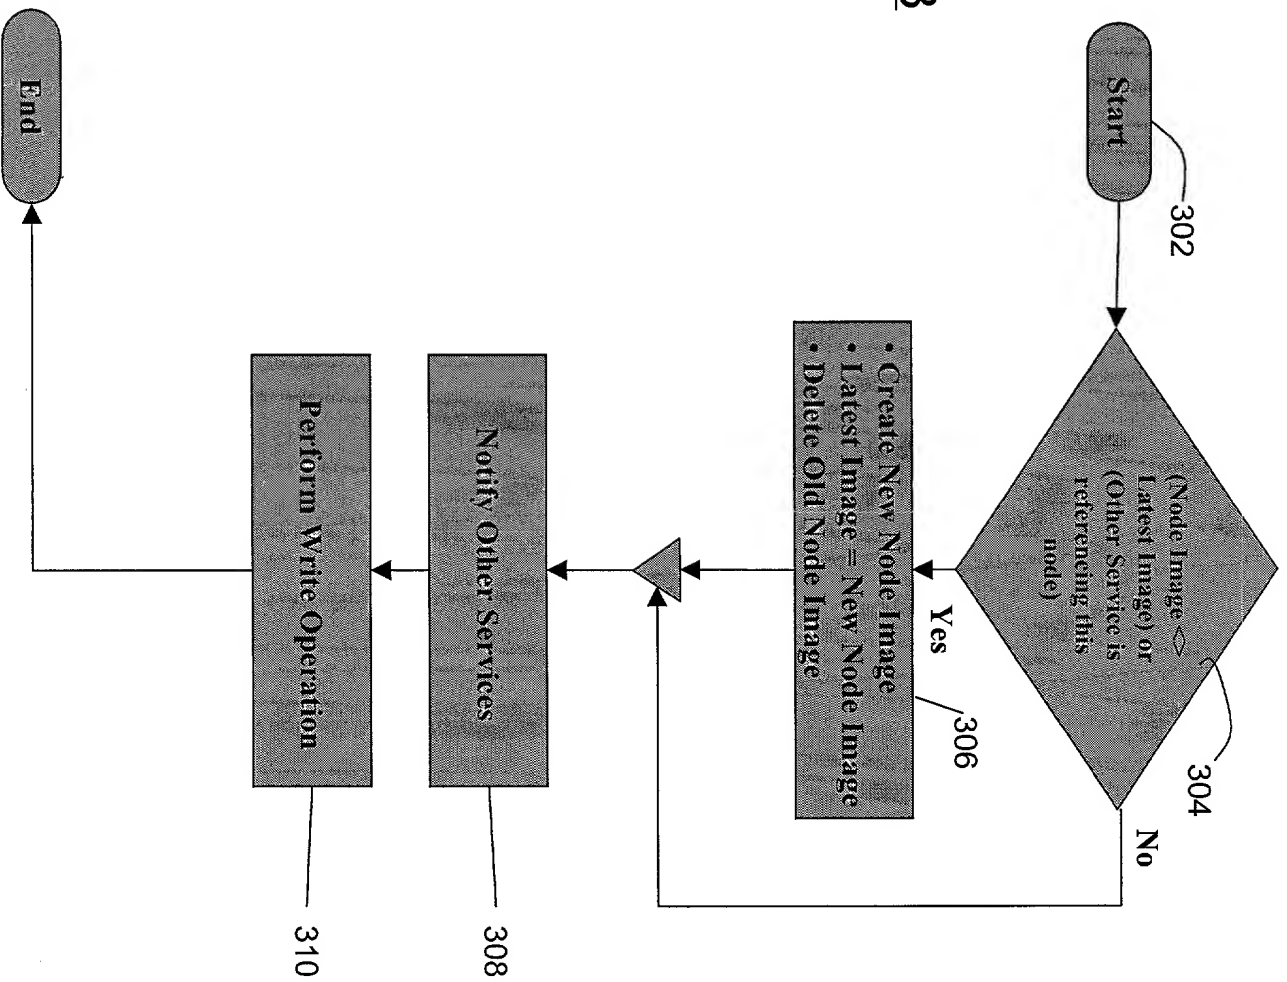


Fig. 3







Fig. 4

Attributes are read by Services, Attributes are written indirectly by creating write operations which are performed by a single unique Service, the Simulation Service. The process of updating attributes is "frame safe". Each Service can be running at a different "frame rate".



Services are running at independent frame rates (using different periodic time).
Each Service is running in a separate thread

Fig. 6

